

SPECIFICATION

Electronic Version 1.2.8

Stylesheet Version 1.0

IN-LINE CODE SUPPRESSION

Background of the Invention

- [0001] Field of the Invention The present invention generally relates to the operation of digitally controlled electronic devices such as digital data processors, digital processor controlled "smart" devices and the like operable in a plurality of modes and, more particularly, reduction of software code overhead based on operational mode.
- [0002] Description of the Prior Art Portable personal computers and many other conveniences that include digital processors (e.g. processor controlled measurement devices, smart pagers, "palm-top" devices and the like) are widely available and used in many different environments. Demands for increased functionality and sophisticated graphic user interfaces have required applications programs used in such devices which can be very large. Further, it is common for a plurality of applications programs to be available concurrently on such devices.
- [0003] At the present time, the cost of software can be a major portion of the cost of the device and constitute a major portion of the processor overhead, limiting performance and often greatly extending initialization time when such a device is started. Initialization procedures generally include extended testing for operational mode parameters and connected devices which may or may not be supported by the application. Results of this testing are generally stored in a configuration table that is interrogated each time particular functions of the application are run and generally result in a plurality of conditional branches to perform the function; thus consuming substantial processing time during both initialization and execution of an application program. This amount of processing time may be particularly critical when it may become necessary to re-start the processor and certainly prevents re-starting from being accomplished in real time.

[0004] Further, such overhead may not be fully consistent with desired operational environments and modes. For example, during operation of a portable personal computer on battery power, it is desirable to minimize power consumption. However, the processing overhead for maintaining availability of all software may be less desirable than extending the period of possible operation on battery power, particularly since a full complement of peripherals are much less likely to be connected when, for example, a lap-top computer is used in an environment where battery power is necessary.

[0005] Thus, it may be desirable to reduce processor overhead by limiting the amount of code decoded and/or stored for possible use. For example, code such as routines involving complex graphic user interfaces, printer drivers when no printer is connected, communications arrangements when no network access is available and the like, which are not accessed in certain operating modes (but which must be included in the applications to support the full extent of desired functionality in other operating modes and environments where, for example, power consumption is of lower importance) may be suppressed and effectively removed from the application. Conversely, when a device is connected to normal alternating current power sources, it may be desirable to disable certain power-saving features such as display blanking that may be inconvenient to a user.

[0006] At the present time, such code can exceed well in excess of 10% of the instruction stream of an application. If the application is to support many different configurations, such codes that may be desirable to disable are usually distributed in the application with an in-line granularity that makes it difficult to turn on or turn off individual operations or sections of code. Current operating systems impose the constraint of maintaining the complete code, including all possibly needed conditional branch trees (sometimes referred to as code cones) of an application for testing for operational mode parameters and identifying attached peripheral devices and then determine what functions are to be enabled. This circumstance is aggravated by the increasing number of ports which may be made available for connection of peripheral devices.

[0007]

Due to the nature of the start-up code a large processing overhead and power drain may be imposed when an application is started when only small sections of code or single operations (opcodes) may require suppression, although the number of opcodes and small

sections of code to be suppressed may be very numerous. Therefore, even configuring the system to limit power consumption imposes a significant power requirement. More specifically, at start-up, processors need to determine which functions are supported or have been changed. Usually, the processor follows a series of branch trees to test either memory locations or device status bits and then allow a section of code to be executed or skipped during the current power-on machine cycles depending on the results of the tests. Each time the code enters a turned-off domain, re-running of a test may be required and the code branched around at the cost of execution time and power.

[0008] Moreover, the amount of time required for this processing is significant and real-time re-starting of digital processors is not currently possible. Reduction of re-start time has generally been approached by limiting the amount of testing performed on the processor and associated memory and thus may not ensure reliable operation. Nevertheless, many applications have been identified such as machine or vehicle control where re-start time may be extremely critical.

Brief Summary of the Invention

[0009] It is therefore an object of the present invention to provide a technique for suppression of single opcodes or complete blocks of code that can be skipped in order to disable their function and eliminate processing overhead associated therewith and thus requiring only minimal processor overhead.

[0010] It is another object of the invention to provide a technique of improving processor performance, such as reduced power consumption or increased speed, by suppressing opcodes and sections of code which may limit desired performance relative to some operating environments or which are not related to functions supported by a given configuration.

[0011] It is a further object of the invention to provide particular functions to be selectively enabled or disabled without deletion of instruction code from application programs.

[0012] It is yet another object of the invention to reduce processor start-up and re-start time without reducing functional testing of the processor, memory and associated peripheral devices.

[0013] In order to accomplish these and other objects of the invention, a method of operating a digital system controlled by operation codes and operable in a plurality of operational modes is provided comprising steps of processing an application program to insert execution bits in operational codes preceding instructions which are not used in particular operational modes, evaluating each operational code prior to decoding of said each operational code, skipping an operational code responsive to detecting a particular state of an execution bit in a preceding instruction, and decoding remaining operational codes.

[0014] In accordance with another aspect of the invention, a processor is provided comprising a sequence of instructions, each said instruction including an execution bit, and means for bypassing an instruction of said sequence based on a particular state of an execution bit in a current instruction.

Brief Description of the Several Views of the Drawings

[0015] The foregoing and other objects, aspects and advantages will be better understood from the following detailed description of a preferred embodiment of the invention with reference to the drawings, in which:

[0016] Figure 1 is a simplified block diagram illustrating operation of the invention,

[0017] Figure 2 is a schematic illustration of the operation of a simple implementation of the invention,

[0018] Figure 3 is a schematic illustration of a second implementation or alternative function to that of Figure 2 in accordance with the invention, and

[0019] Figure 4 is a schematic illustration of a third implementation or alternative function to that of Figure 2 in accordance with the invention.

Detailed Description of the Invention

[0020] Referring now to the drawings, and more particularly to Figure 1, there is shown a high level block diagram illustrating operation of the invention. The invention begins with current code 100 of an application which must be fetched from storage prior to being executed by a processor. Normally lines or blocks of code are fetched from a mass memory, decoded and stored into a smaller electronic memory from which they may be

obtained as needed in less time than that required for a fetch from mass memory.

[0021] The invention is preferably implemented by post-processing (e.g. processing after the application is built and compiled) program code 100 as shown at 110 to set an execution bit associated with the operation code (opcode) immediately prior to the opcode to be suppressed. This processing of the application code can be done in many ways, the details of which are not critical to the practice of the invention. However, it is preferred to perform the processing based on a previous operating state of the processor which may be input as various criteria 160 to determine whether a particular opcode or section of code will be used.

[0022] When this execution bit is active, the next opcode is regarded as a "don't care", as detected by pre-fetch evaluation 120. Thus, when the opcode with the active execution bit active reaches the pre-decode stages the program counter is incremented (140) over the next instruction 130 to instruction 150 and instruction 130 is suppressed and not decoded or otherwise used by the processor or stored for processor use; thus eliminating processor overhead associated therewith and even the decoding overhead associated with the suppressed codes.

[0023] This mode of operation is illustrated in greater detail in Figure 2. Standard instructions 200 and 210 have their execution bits 205, 215 set at "0" (inactive) and are handled and decoded normally and in sequence. Standard operation 220 is also handled normally but since the execution bit 225 is set to "1" (active), low power instruction 230 is skipped/bypassed (regardless of the state of its execution bit) and decoding proceeds with standard instruction 240 and following instructions in sequence. If the execution bit of standard instruction 240 was previously set to "1", however, the next instruction following instruction 240 would be skipped/bypassed.

[0024] In essence, this process makes a conditional branch available at every opcode so that selected lines of code can be suppressed easily, simply and flexibly with virtually no processing or decoding overhead. This effect of the invention can also be viewed as substituting the conditional branch provided by the invention for one or more (and potentially many) conditional branches executed each time particular functions are required based on operating mode and configuration data collected during initialization. That is, for

each operating mode, function or supported device, a single conditional branch is made to suppress code prior to decoding and running the application program and thus avoid potentially many more time-consuming conditional branch instructions that may be repetitively executed during initialization and execution of the application program.

[0025] The process of suppressing a single opcode discussed above in regard to Figure 2 can be extended to blocks of data in several ways; preferred, exemplary techniques being illustrated in Figures 3 and 4. For example in Figure 3 setting of the execution bit of instruction 300 causes all following instructions 310, 320, 330 to be skipped until another active execution bit is encountered in instruction 340. This control can be very simply exercised by toggling a status bit in a register in accordance with active execution bits of respective instructions and determining whether skipping is to be performed by comparison of the status bit with the execution bit. In this manner, instructions in blocks can be iteratively skipped in a simple manner using only a single execution bit per instruction.

[0026] Alternatively, as shown in Figure 4, more than one execution bit is provided and may be used to store an offset address. When such a group of execution bits is encountered in, for example, instruction 400, the program counter is incremented by that number to provide the address/sequential location of the next instruction to be decoded. Thus, the process can branch directly to instruction 420 while skipping block 410. This process is faster than that of Figure 3, discussed above, since the skipping/conditional branch process is direct rather than iterative but requires provision and setting of additional execution bits.

[0027]

It should be appreciated that the skipping of instructions has no effect on accessibility of remaining instructions from a memory local to the processor such as a working memory. All instructions remain in correct sequence in the application program and are accessible from main memory such as a mass storage in the usual manner. As the instructions are evaluated prior to decoding the skipped instructions are not decoded and the remaining instructions are decoded and the result stored in local memory while address translation data is accumulated in a look-up table to translate the addresses of the original program sequence to the address in local memory, as is known in the art. However, a meritorious

effect of applying the invention in such a known environment is that the memory local to the processor, such as a cache, is not necessarily burdened with storing the skipped instructions and thus can be effectively expanded in capacity; increasing cache/local memory hit rates and processor performance. However, it is often preferable to store the skipped instructions in the program stream but simply to suppress decoding in order to avoid the processing overhead associated therewith.

[0028] In view of the foregoing, it is seen that the invention provides a simple and easily implemented technique of improving processor performance by reducing instruction code maintenance overhead with little, if any, burden on the processor from its implementation. The reduction in processing overhead is sufficient to support the function of real-time processor restarts since the effects of the invention on processor performance are particularly evident during initialization operations because no deactivation of individual opcodes and code segments is required because they have been effectively removed from the code stream before reaching the processor. Power savings are also achieved since no processing is required to maintain opcodes or code segments which are not supported by a particular operating environment or mode.

[0029] While the invention has been described in terms of a single preferred embodiment, those skilled in the art will recognize that the invention can be practiced with modification within the spirit and scope of the appended claims.